arXiv:1608.04819v2 [math.NA] 20 Mar 2017

Parameter Selection for HOTV Regularization

Toby Sanders
School of Mathematical and Statistical Sciences, Arizona State University, Tempe, AZ, USA.

Abstract

Popular methods for finding regularized solutions to inverse problems include sparsity promoting ℓ_1 regularization techniques, one in particular which is the well known total variation (TV) regularization. More recently, several higher order (HO) methods similar to TV have been proposed, which we generally refer to as HOTV methods. In this letter, we investigate the problem of the often debated selection of λ, the parameter used to carefully balance the interplay between data fitting and regularization terms. We theoretically for a scaling of the operators for a uniform parameter selection for all orders of HOTV regularization. In particular, parameter selection for all orders of HOTV may be determined by scaling an initial parameter for TV, which the imaging community may be more familiar with. We also provide several numerical results which justify our theoretical findings.

Keywords

Image reconstruction, Regularization, Inverse problems, Higher order TV, Parameter selection

1 Introduction

Inverse and imaging problems play a vital role in practically all areas of technology and science. These problems naturally turn up in applications[1], whenever data measurements are collected that must be processed or *inverted* to form resulting images that explain the measurements. A popular formulation for inverse signal and imaging problems takes the form

$$\min_f F(f) + H(f), \tag{1}$$

where f is the image we want to recover, F is a data fitting term, and H is a regularization term that encourages some prescribed smoothness on the image.

If b is the known noisy data, and an estimated forward operator A is known that maps f to the data, then the data fitting term is commonly of the form

$$F(f) = \frac{\lambda}{2}\|A(f) - b\|_2^2. \tag{2}$$

Typically, A is a linear operator that can be approximated by a matrix, and $\lambda > 0$ is a model parameter which is chosen to achieve the best results by properly balancing F and H.

The regularization term H is usually a convex norm or semi-norm. Historically H has been selected as the ℓ_2 energy norm of the image to improve ill-conditioning of some inverse problems, a famous method which goes by the name of Tikhonov regularization. Our interests in this paper are more recently popular

[1]MRI, X-ray tomography, optical coherence tomography, and synthetic aperture radar, just to scratch the surface.

ℓ_1 regularization methods which are often categorized as compressed sensing [7]. These methods promote sparsity of f in some selected domain by setting the regularization term to

$$H(f) = \|Tf\|_1, \tag{3}$$

where T is the mapping under which the solution is assumed to be sparse. A very common regularization choice for imaging problems is to use the TV norm, which is equivalent to setting T as the first order finite difference operator that maps f to the difference between all adjacent pixels of f. This type of regularization encourages sparsity in the boundaries or number of jumps in the image [16, 22]. Also commonly used are wavelet frames and bases [14, 20], and more recently higher order finite difference transforms, an approach sometimes referred to as higher order TV [8, 1]. We now address the main focus of this paper, the famous and often debated problem of the selection of the parameter λ in (2), particularly for higher order TV methods.

With our formulation, it is most generally accepted that the parameter λ is chosen to be smaller for noisier data b, and chosen to be larger for more accurate data values. With regards to TV regularization, numerous publications have already been devoted to parameter selection [25, 3, 23], and many working in the imaging community may be familiar with this task. Indeed, the motivation behind this work was to generalize the parameter selection based off of TV. Having such a parameter selection may prove to be useful in many practical applications. With general image reconstruction problems, it may not be known exactly which regularization method is appropriate prior to searching for a solution, i.e. precise characterization of the underlying image structure may not be known prior to reconstruction. Therefore several regularization approaches may be carried out in conjunction to realize the most effective approach [19]. Having an accurate parameter scaling between the approaches is necessary for an unbiased comparison of the various approaches and limits any additional parameter tuning.

The referenced approaches for TV parameter selection resolve to iterative selection of λ by posterior estimates, or by a priori estimation of noise levels. Another alternative for general regularization problems is the popular L-curve method [10, 6]. This method is essentially a brute force approach, in which resulting values for F and H in (1) are plotted together for an array of parameter selections in order to determine the parameter value that yields the optimal payoff between the F and H. This approach has only been formally studied for Tikhonov regularization, and more importantly it requires one compute many solutions to the problem.

In this paper, we work towards uniform parameter selection for the more general ℓ_1 methods that we refer to as higher order total variation (HOTV). A number of HOTV type of approaches have been proposed [8, 5, 15, 11]. While each of these methods have subtle differences, the motivation behind each is somewhat similar: to allow for more general function behavior and reduce staircasing effects observed with TV by removing penalties for low order polynomials. This order is usually flexible within the approach, and the removal of such penalties may be achieved by defining the transform T to in some way approximate higher order derivatives. For instance, in [5] the authors design a regularization term in the continuous setting based on higher order weak derivates and show that low order polynomials are not penalized, again depending on the order of the method. In [8], a fourth order PDE-based technique was used to a similar effect, also removing penalties for low order polynomials. In [15] the authors used within the regularization terms both first order derivatives with second order derivates and demonstrated reduced staircasing effects. The algorithmic design in each of these approaches are somewhat different depending on the date of publication, but at the current time the alternating direction method of multipliers approach has primarily been adopted [13, 24, 4].

We work in particular with the HOTV formulation that has been referred to as polynomial annihilation (PA) regularization [1, 2], which takes a straight forward approach to removing penalties from low order polynomials. This is achieved by using the common discretized approximation of derivatives otherwise known as finite differences. The MATLAB algorithm for this approach can be downloaded with

a corresponding user's guide online [17]. Compared with other HOTV approaches, which have primarily been demonstrated for denoising and deconvolution problems, our algorithm has been shown effective for a variety of 2D and 3D inverse problems. In [19] it was used to reconstruct 3D nanoscale volumes with dimensions 2^{28} total voxels from undersampled tomographic data of total dimension approximately 2^{26}. In [18] the approach was adapted for a joint sparsity model to reconstruct scenes from nonuniform Fourier related synthetic aperture radar data sets.

The ability to effectively solve such large and complex imaging problems is facilitated by a robust algorithm with simple parameter selections. This work details some important results that has enabled the ease of use of these algorithms, and our numerical studies here additionally confirm the use of these results. We first argue analytically in two distinct approaches that precise scaling of λ depends on the order of the HOTV method used, and the result essentially tells us that we should scale by powers of two[2]. While a simple scaling by two may seem elementary, the details that arrive at this result are nontrivial, and the approach has been extremely effective [19]. We provide robust numerical results that confirm our choice of the scaling of λ, by both imaging examples and a large series of simulations where we precisely determine the optimal parameters. In section 2 we introduce our HOTV regularization method. In section 3 we provide our analysis for appropriate parameter selection, and in section 4 we present several numerical experiments that confirm our proposed parameter selection.

2 Polynomial Annihilation for ℓ_1 Regularization

The linear transform T in (3) has more recently gained attention as a higher order finite difference operator, and in particular one known as polynomial annihilation (PA) regularization [1, 2, 21]. Loosely speaking, the transform for PA annihilates f at grid points where f is essentially a polynomial of degree less than a desired value, which is called the order of the PA operator. We denote the PA operator of order k by T_k, while noting that T_1 is equivalent to TV. As with TV, the PA regularization encourages sparsity in the number jumps in the image. However, in contrast with TV, characterization of the smooth regions is relaxed, where they are modeled as low order polynomials (e.g. degree 2 or 3) rather than constant functions.

Without formal details, we state that on an equally spaced grid the PA operator is equivalent to a kth order finite difference. Hence, the PA transform has the following definition:

Definition 1. *Let $f \in \mathbb{R}^N$ and let $j \leq N-k$. Then the j^{th} element of the order k PA transform of f is given by*

$$(T_k f)_j = \sum_{m=0}^{k} (-1)^{k+m} \binom{k}{m} f_{j+m}. \tag{4}$$

For indices $j > N-k$, we may also define $(T_k f)_j$ using a periodic extension of f.

For higher order finite differences in multiple dimensions, we simply extend T_k to evaluate the difference along each dimension. For instance, in 2D, we may have $T_k = [T_k^x, T_k^y]^T$, where T_k^x computes the differences along the x dimension and similarly for y. For example, for $f \in R^{n\times n}$,

$$(T_k^x f)_{j,\ell} = \sum_{m=0}^{k} (-1)^{k+m} \binom{k}{m} f_{j+m,\ell}.$$

To briefly see that T_k defined in (4) indeed annihilates polynomials of degree less than k, one can first consider a function defined on a continuous domain. Here, the operator is evidently the k^{th} derivative

[2]Although in this work we argue for a scaling of the parameter, the approach has been so successful that the scaling of the parameter has now been removed as an option in the algorithm, and the operators are scaled instead. Therefore the user may simply choose parameter values without concern for the order of the method.

operator. Then discretization of f along with the mean value theorem leads one to find that standard finite differences yield the discretized version of the PA operator, which is well known to be (4) (see for instance, [12]).

Finally, the PA regularization method can be expressed as finding the solution to

$$\min_f \Big\{ \frac{\lambda}{2} \|Af - b\|_2^2 + \|T_k f\|_1 \Big\}, \tag{5}$$

where T_k is defined by (4).

In order to solve (5) for large imaging problems, an efficient iterative algorithm is needed. In our work, we have followed the alternating direction method of multipliers (ADMM) approach in [13], which rewrites the problem with an augmented Lagrangian function and uses alternating direction minimization to solve equivalent subproblems. Our algorithm is available online [17].

3 Selection of λ

In this section we address the appropriate selection of the parameter λ in (5), showing explicitly how the parameter depends on the order k of T_k. In general, when solving a problem such as (5), appropriate selection of λ is greatly eased by rescaling the operators A and T_k, in order to reduce the dependence on the specific problem and regularization transform. In doing so, λ is then primarily dependent on the noise levels and error in the forward model. [3]

To reduce problem dependence, A is rescaled so that $\|A\|_2 = 1$, and the data values are rescaled accordingly. Determining this rescaling factor can be accomplished relatively easily, since for $A \in \mathbb{R}^{m\times n}$, $\|A\|_2^2 = \rho(A^T A)$, the largest eigenvalue of $A^T A$ (see, for instance [9]). This eigenvalue can be computed iteratively in computational time that is insignificant to the time needed to solve (5), making this rescaling a practical rule.

Similarly, we may rescale T_k so that $\|T_k\|_1 = 1$. However, for simplicity we keep the integer values defined in the operator given by (4), and thus instead choose to rescale λ depending on $\|T_k\|_1$. To do this, we start with the following proposition:

Proposition 1. *For the operator T_k defined in (4), the ℓ_1 norm of T_k is given by*

$$\|T_k\|_1 = 2^k. \tag{6}$$

While the exact details of the proof are left to the reader, the general idea is outlined below. First recall that the ℓ_1 norm of a matrix is equivalent to the maximum of the ℓ_1 norm of the columns, [9]. Because the operator is circulant [4], the sum of absolute values of each column of T_k are equivalent, leading to

$$||T_k||_1 = \max_{1\le j\le n} \sum_{i=1}^{N} |(T_k)_{ij}| = \sum_{i=0}^{k} \left| (-1)^{k+i} \binom{k}{i} \right| = 2^k. \tag{7}$$

Thus by definition of the matrix norm (see again, [9]) we have

$$||T_k f||_1 \le 2^k ||f||_1, \tag{8}$$

[3] We note that it is important to make the distinction between noise levels and model errors, since the latter is typically coherent and cannot be captured by standard noise models. Furthermore, in some cases this means $\|Af - b\|_2$ can not be made arbitrarily small, and thus λ should be chosen accordingly.

[4] In our definition, the operator is only truly circulant if the periodic extension is used. Nevertheless, the resulting norm is the same in both cases.

suggesting that the regularization parameter in (5) using T_k be of the form

$$\lambda = 2^{k-1}\lambda_1, \tag{9}$$

for some appropriate choice of λ_1 for TV. We emphasis again here that numerous publications have already been devoted to parameter selection for TV [25, 3, 23], making a scaling factor for HOTV based off a selection for TV a reasonable approach. A precise scaling is also useful if one wants to make unbiased comparisons between any two of the approaches.

It is important to note that equality may only be achieved in (8) for signals that are not typically of interest. For example, $\mathbf{f} = \{(-1)^j\}_{j=1}^N$ represents such a signal, but is of limited practical use. Indeed we are interested in a smaller class of *smooth* signals for which equality may not hold. Nevertheless, as is shown in what follows, choosing λ using (9) is robust and removes the need for additional parameter tuning for each k used.

The following is our main result, that along with (6) also justifies our scaling.

Theorem 1. *Let f be a piecewise constant signal in $\mathbb{R}^N$ with jump locations at the indices over the set defined by $S = \{j \mid f_{j+1} - f_j \neq 0\}$. If the jump locations are sufficiently spaced so that for any $j_1, j_2 \in S$, we have $|j_1 - j_2| > k$, then the following relation between the TV and PA norms holds:*

$$\|T_k f\|_1 = 2^{k-1}\|T_1 f\|_1. \tag{10}$$

Theorem 1 is consistent with (9) for choosing λ in (5), and this scaling is further justified by our numerical results. Although the statement is for piecewise constant functions, in general the main contribution of to $\|T_k f\|_1$ is at the jumps, even when the function is not a piecewise constant. We believe a similar argument can be made for higher order piecewise polynomials, but this is notably more difficult and is reserved for potentially later work.

We now prove Theorem 1, which will follow almost immediately by the Lemma 1 below. We will first need the following proposition:

Proposition 2. *For all positive integers k and m with $k > m$, we have*

$$\binom{k-1}{m} = \sum_{j=0}^{m} \binom{k}{j}(-1)^{j+m}. \tag{11}$$

This identity can be proven with repeated application of the identity $\binom{k-1}{m} = \binom{k}{m} - \binom{k-1}{m-1}$, starting with the left hand side until the complete alternating summation arises.

To prove Theorem 1, we only need to consider a simple example from the class of general piecewise constant functions defined in the theorem. Let the smooth signal f be defined by

$$f_j = \begin{cases} \alpha & \text{if } j \leq i \\ \beta & \text{if } j > i. \end{cases} \tag{12}$$

Here $\alpha \neq \beta$ are constants, and i is some arbitrary index sufficiently far from the boundary. The following lemma illustrates that $\|T_k f\|_1$ for f in (12) increases by a factor of 2 for each k.

Lemma 1. *The discrete signal given by (12) satisfies*

$$\|T_k f\|_1 = 2^{k-1}|\alpha - \beta|. \tag{13}$$

Proof. First observe the evaluation of $(T_k f)_m$ vanishes at points away from the index i. In particular, $(T_k f)_m = 0$ for $m = 1, 2, \ldots, i-k, i+1, i+2, \ldots, N-k$.

However, at index $m = i - \ell$, with $0 \leq \ell < k$, we have

$$\begin{aligned}(T_k f)_m &= \alpha \left(\sum_{j=0}^{\ell} (-1)^{j+k} \binom{k}{j} \right) + \beta \left(\sum_{j=\ell+1}^{k} (-1)^{j+k} \binom{k}{j} \right) \\ &= (\alpha - \beta) \sum_{j=0}^{\ell} (-1)^{j+k} \binom{k}{j},\end{aligned}$$

where we have used the fact that

$$\sum_{j=0}^{\ell} (-1)^{j+k} \binom{k}{j} = - \sum_{j=\ell+1}^{k} (-1)^{j+k} \binom{k}{j}.$$

Applying (11) yields $(T_k f)_m = (\alpha - \beta)\binom{k-1}{\ell}(-1)^{k-\ell}$. Finally, taking the ℓ_1 norm we obtain

$$\|T_k f\|_1 = |\alpha - \beta| \sum_{\ell=0}^{k-1} \binom{k-1}{\ell} = 2^{k-1} |\alpha - \beta|. \tag{14}$$

□

Theorem 1 now follows immediately from Lemma 1 by generalizing this discrete signal in (12) to a piecewise constant signal with sufficiently spaced jump locations. We also note that this result generalizes without any trouble to two and three-dimensional signals.

4 Numerical Experiments

4.1 Numerical Determination of Optimal λ

To determine the exactness of our scaling in a robust numerical fashion we set up 1,000 1-D numerical simulations, in which the optimal λ is determined for each simulation. The simulations were each uniquely set up in the following way:

- **Signal construction:** In each simulation, a piecewise polynomial signal was randomly generated over a 256 point mesh. Each signal was generated by first randomly designating some number of jumps for the signal, between 2 and 20 jumps with each possible number having equal probability. Then the location for each jump was randomly determined. Over each continuous section between the jumps, a polynomial of either order 0, 1, or 2 is generated for the each section, where again this order is randomly chosen and each order has equal probability.

- **Sampling:** A random sampling matrix A was generated to sample the generated signal, where the sampling rate is randomly chosen between 25% and 100%. The sampling rate is defined by the number of samples, or rows of A, divided by the number of grid points, or columns of A. The entries of A are determined by randomly choosing 10% of the entries to be nonzero, where the nonzero entries are uniformly distributed over the interval $[0, 1]$.

- **Noise:** Mean zero Gaussian white noise was added to the data samples, where the standard deviation of the noise is set to be a random value between 0 and 3, each value having equal probability of occurring.

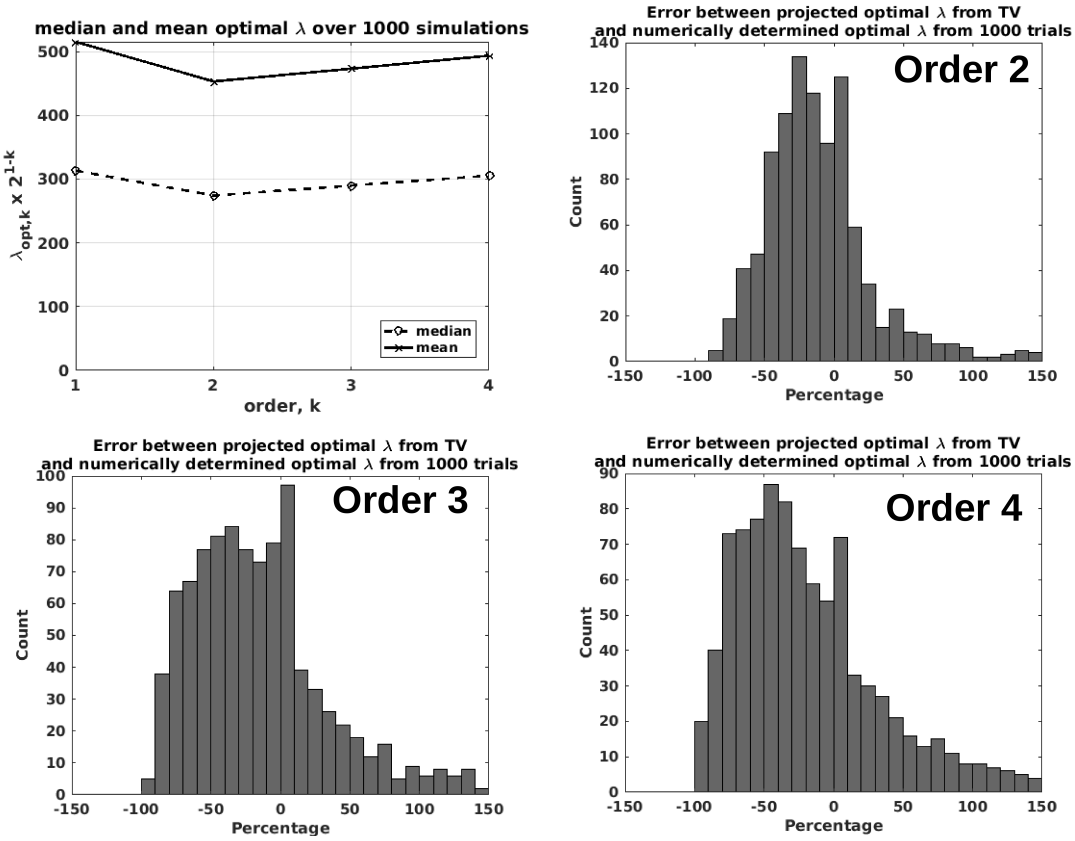

Figure 1: Results for optimal λ determined from 1000 simulations. For each order (2,3,4), the histograms show the errors between the projected optimal λ from TV as given by (9) and the numerically determined optimal λ determined by (15). Shown in the top left panel are the mean and median optimal λ values from all 1,000 simulations for each order (the values are scaled by 2^{1-k} for each order k).

This set up gives us an extremely wide variety of examples to draw conclusions and eliminates any possible biases in the particular set up.

Once the problem is set up, the optimal λ is determined by a brute force search, which is terminated once the optimal value is determined within a sufficient tolerance. We choose the optimal λ to be defined as

$$\lambda_{opt,k} = \arg\min_{\lambda} \left\{ ||f_{true} - f^*||_2 \;\middle|\; f^* = \arg\min_{f} \frac{\lambda}{2}\|Af - b\|_2^2 + \|T_k f\|_1 \right\}, \tag{15}$$

where f_{true} is the original signal. In other words, the optimal λ is that which yields a solution closest to the original function.

According to our work, ideally the optimal lambda for order k is given by $\lambda_{opt,k} = 2^{k-1} \cdot \lambda_{opt,1}$. Therefore we define the relative error of our approach by

$$\frac{2^{1-k} \cdot \lambda_{opt,k} - \lambda_{opt,1}}{\lambda_{opt,1}}, \tag{16}$$

which is calculated for all 1000 simulations and for $k = 2, 3, 4$.

The results from the simulations are presented in Figure 1. For each order, the frequency of each relative error (as a percentage) is shown in histogram plot. Each shows a small offset centered curve to the left of 0%, while also showing a secondary peak directly at 0%. There is a skew to the right however. This pulls the overall average optimal λ to the right of the main peaks. The mean and median optimal λ is plotted in the top left panel, where for each order the values are rescaled by 2^{1-k}, so that by our scaling of λ we would ideally see straight line for these two plots. There is a modest dip for orders 2 and 3, but not significantly so. Generally speaking these results confirm that our scaling yields fairly accurate results. However, the histogram plots show wider curves for higher orders, suggesting that the scaling is more often accurate for smaller orders and less reliable for larger orders.

4.2 Two-dimensional Imaging Experiments

For two-dimensional numerical experiments we reconstruct two phantom images of size 256×256 from undersampled and noisy data. The phantoms are shown in Figures 2 and 3 (a). The first phantom is the classical Shepp-Logan phantom, which is piecewise continuous and well suited for TV. The second phantom is similar to a cross-sectional brain image and has many nontrivial image features making it better suited for HOTV methods.

The sampling of the images was at a rate of 50%, giving $256^2/2$ total data points, which are corrupted with Gaussian white noise yielding an SNR of 23.75. The sampling matrices are generated just as in the previous 1-D set up.

Reconstructions from the data were performed with HOTV regularization of order 1 (TV), 2, 3, and 4, both with and without the suggested scaling of λ based off an initial selection for the TV reconstruction. Some of the resulting reconstructed images are shown in Figure 1 and 2 (b)-(f). The least squares solution in (b) is shown for baseline comparison and to illustrate the difficulty of reconstructing in this setting. For the Shepp-Logan phantom, some of the fine features present in (c-e) with the scaling of λ become absent in (f) where the upscaling of λ is not used. Particularly the three small ellipses at the bottom of each become difficult to make out when the scaling is not used.

For the second phantom, the TV model smooths over the fine details of the image. A comparison with the order 3 approach with the scaling (e) and without the scaling (f) are again made. These again indicate that λ is too small in (f), and some of the accurate image features that are seen in (e) are smoothed together in (f).

Perhaps most revealing are the plots in Figure 1 (g) and (h). In (g), the corresponding PA ℓ_1 norms, $\|T_k f\|_1$, for each reconstruction are computed and compared with each corresponding norm for the phantom

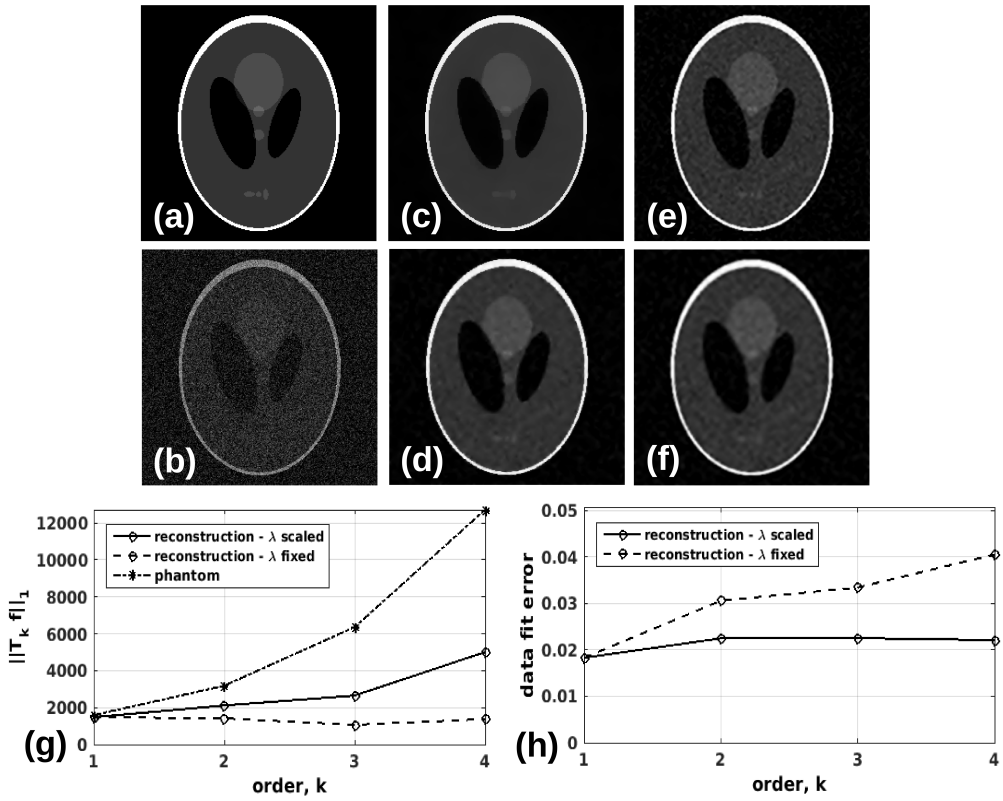

Figure 2: (a) Shepp-Logan phantom. (b) Least squares reconstruction. (c) TV reconstruction. (d-e) HOTV reconstructions of orders 2 and 3 respectively with λ scaled according to (9). (f) HOTV order 3 reconstruction without scaling λ. (g) HOTV norms for each reconstructions compared with the phantom. (h) Relative data fitting error for each reconstruction.

image. First note that the norms for the phantom closely follow our theoretical result in Theorem 1 of an increase by a factor of two for each increase in the order. Without the scaling of λ, the resulting norms of the reconstructions essentially remain flat, where we already observed the true norm of the phantom increases by a factor of around two. Thus with the scaling of λ, the reconstructions follow the trend for the phantom. In (h), the relative error of the reconstructions is shown, which is defined by $\|Af-b\|_2/\|b\|_2$. It is clearly shown that for reconstruction without scaling λ, the relative error increases significantly for each increase in the order. This indicates that λ is too small for these higher order methods. On the other hand, the reconstructions with the scaled λ show comparable relative errors, as one would hope for.

5 Conclusions

We investigated the common problem of parameter selection for regularized inverse problems, in particular for our successful TV and HOTV methods known as polynomial annihilation (PA). Our analysis given by Proposition 1 and Theorem 1 showed that for each increase in the order of the method, the parameter λ should be increased by a factor of two. These theoretical results essentially summarize that for a certain class of signals, the ℓ_1 norm of our HOTV transform of these signals increases by a factor of two for each increase in the order, and thus λ should be scaled according to (9). Two-dimensional numerical results justified this scaling of λ, particularly showing that when the scaling was carried out, the data fitting error was nearly even for all orders, and the increase in the ℓ_1 norms trended with the true phantom norms

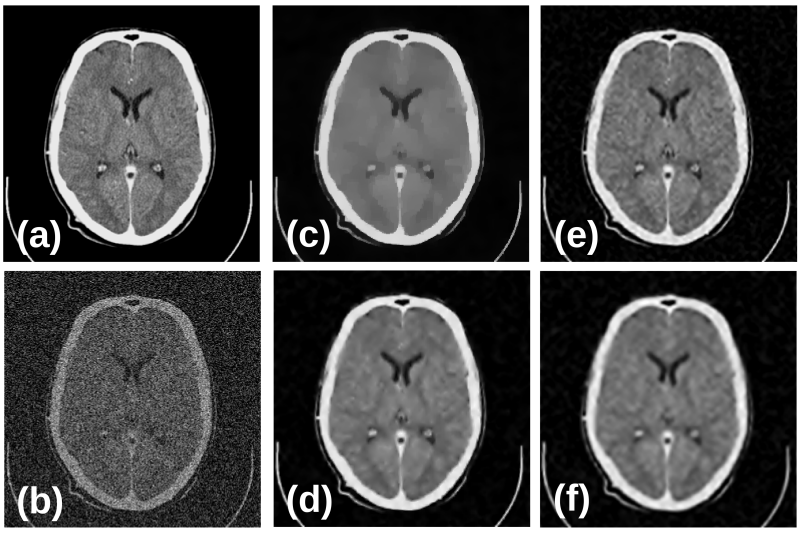

Figure 3: (a) Phantom image. (b) Least squares reconstruction. (c) TV reconstruction. (d-e) HOTV reconstructions of orders 2 and 3 respectively with λ scaled according to (9). (f) HOTV order 3 reconstruction without scaling λ.

that agreed with the theoretical results. More robust numerical analysis was carried out in a set of 1,000 1-D simulations, and showed that the optimal λ for each order generally followed the trend of scaling according to (9) with some modest error. This scaling allows for unbiased comparisons between the varied approaches and may be especially useful for anyone already familiar with the classical TV approach.

Acknowledgements

This work is supported in part by the grants NSF-DMS 1502640 and AFOSR FA9550-15-1-0152.

References

[1] R. Archibald, A. Gelb, and R. B. Platte. Image reconstruction from undersampled Fourier data using the polynomial annihilation transform. *J. Sci. Comput.*, pages 1–21, 2015.

[2] R. Archibald, A. Gelb, and J. Yoon. Polynomial fitting for edge detection in irregularly sampled signals and images. *SIAM J. Numer. Anal.*, 43(1):259–279, Jan. 2005.

[3] J.-F. Aujol, G. Gilboa, T. Chan, and S. Osher. Structure-texture image decomposition—modeling, algorithms, and parameter selection. *International Journal of Computer Vision*, 67(1):111–136, 2006.

[4] S. Boyd. Alternating direction method of multipliers. In *Talk at NIPS Workshop on Optimization and Machine Learning*, 2011.

[5] K. Bredies, K. Kunisch, and T. Pock. Total generalized variation. *SIAM J. Imaging Sci.*, 3(3):492–526, 2010.

[6] D. Calvetti, S. Morigi, L. Reichel, and F. Sgallari. Tikhonov regularization and the l-curve for large discrete ill-posed problems. *Journal of computational and applied mathematics*, 123(1):423–446, 2000.

[7] E. Candès and J. Romberg. Sparsity and incoherence in compressive sampling. *Inverse Probl.*, 23(3):969, 2007.

[8] T. Chan, A. Marquina, and P. Mulet. High-order total variation-based image restoration. *SIAM J. Sci. Comput.*, 22(2):503–516, 2000.

[9] G. Golub and C. V. Loan. *Matrix Computations*, pages 54–57. Johns Hopkins University Press, 3rd edition, 1996.

[10] P. C. Hansen. The l-curve and its use in the numerical treatment of inverse problems. In *in Computational Inverse Problems in Electrocardiology, ed. P. Johnston, Advances in Computational Bioengineering*, pages 119–142. WIT Press, 2000.

[11] Y. Hu and M. Jacob. Higher degree total variation (hdtv) regularization for image recovery. *IEEE Transactions on Image Processing*, 21(5):2559–2571, 2012.

[12] R. LeVeque. *Finite Difference Methods for Ordinary and Partial Differential Equations*, chapter Finite Difference Approximations. Society for Industrial and Applied Mathematics, 2007.

[13] C. Li, W. Yin, H. Jiang, and Y. Zhang. An efficient augmented lagrangian method with applications to total variation minimization. *Comput. Optim. Appl.*, 56(3):507–530, 2013.

[14] S. Mallat. *A wavelet tour of signal processing: the sparse way*, pages 292–296. Academic press, 2008.

[15] K. Papafitsoros and C.-B. Schönlieb. A combined first and second order variational approach for image reconstruction. *Journal of mathematical imaging and vision*, 48(2):308–338, 2014.

[16] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1-4):259–268, 1992.

[17] T. Sanders. Matlab imaging algorithms: Image reconstruction, restoration, and alignment, with a focus in tomography. `http://www.toby-sanders.com/software`, `https://doi.org/10.13140/RG.2.2.33492.60801`. Accessed: 2016-19-08.

[18] T. Sanders, A. Gelb, and R. B. Platte. Composite sar imaging using sequential joint sparsity. *Journal of Computational Physics*, 2017.

[19] T. Sanders, A. Gelb, R. B. Platte, I. Arslan, and K. Landskron. Recovering fine details from under-resolved electron tomography data using higher order total variation l1 regularization. *Ultramicroscopy*, 174:97–105, 2017.

[20] J.-L. Starck, F. Murtagh, and J. M. Fadili. *Sparse image and signal processing: wavelets, curvelets, morphological diversity*. Cambridge university press, 2010.

[21] W. Stefan, R. A. Renaut, and A. Gelb. Improved total variation-type regularization using higher order edge detectors. *SIAM J. Imaging Sci.*, 3(2):232–251, 2010.

[22] D. Strong and T. Chan. Edge-preserving and scale-dependent properties of total variation regularization. *Inverse problems*, 19(6):S165, 2003.

[23] D. M. Strong, J.-F. Aujol, and T. F. Chan. Scale recognition, regularization parameter selection, and meyer's g norm in total variation regularization. *Multiscale Modeling & Simulation*, 5(1):273–303, 2006.

[24] S. Tosserams, L. Etman, P. Papalambros, and J. Rooda. An augmented lagrangian relaxation for analytical target cascading using the alternating direction method of multipliers. *Structural and multidisciplinary optimization*, 31(3):176–189, 2006.

[25] Y. W. Wen and R. H. Chan. Parameter selection for total-variation-based image restoration using discrepancy principle. *IEEE Transactions on Image Processing*, 21(4):1770–1781, April 2012.